

[SQUEAKING]

[RUSTLING]

[CLICKING]

JACK HARE: Today, we are going to be discussing neutral particle diagnostics. So the fact of the matter is, we have some plasma, like a tokamak, and it has some block surfaces. And we want to know what's going on inside the core, because this is where the fusion is happening.

In particular, we want to know what the distribution function of the ions is. It's probably Maxwellian-ish, but the "ish" is the interesting bit. There may be some lovely tails of particles here, and those are the ones that might be doing most of our fusion reactions. So we'd like to be able to measure this distribution function f of i .

The trouble is, the particles in the core are hopefully very well-confined. If they're not confined, we haven't done a very good job. So we can't exactly wait for them to wander out and see what they're up to. And so we need some technique of measuring this core distribution function f of i V_i .

Now, we have some ideas based on some of the other diagnostics we've looked at already. For example, we could use electron cyclotron emission. That gave us information about what's going on in the core. But this gives us a measurement of T_e . And T_e does not necessarily have to be the same as T_i . So that's not a very good diagnostic.

We could do something using laser-induced fluorescence, or two-photon absorption laser-induced fluorescence like this. If you remember, that's where we had an upper level and a lower level and we scanned our laser wavelength in frequency space here so that we could address particles which are Doppler shifted towards us and Doppler shifted away from us. And by looking at the intensity of light coming out, we could measure this distribution function here. So this looks like a pretty good technique for measuring the core.

The trouble is that in the core, all of our particles are fully ionized. And so there are no electrons here. So we can't use these techniques because there are no ions with an electron in them in order for us to be able to address it using our laser.

And thirdly, we could also use Thomson scattering, which we haven't talked about yet, but we will do in detail soon. The trouble with Thomson scattering is that it tends to give us information-- scattering-- about the electron distribution function, f_e V_e , like that. And again, that doesn't have to necessarily be the same as the ion distribution function. So we're going to need a new technique in order to measure the distribution of these core particles.

So the idea is, well, what if, spontaneously, some of these core ions turned into neutrals? Well, if some of these core ions turned into neutrals, then they would no longer be confined. They would be able to wander freely out. And perhaps we could stick a detector on the outside of the plasma and measure these neutrals, which will carry with them some information about the ion distribution they came from. And that will allow us to indirectly measure the ion distribution.

So what's a good way to summarize this? They're going to-- these neutrals would give us information on the distribution function. So there's two problems that's coming to your mind straight away. Why exactly would an iron suddenly turn itself into a neutral in the core? And does it stand a chance of actually escaping from the core to the outside, where we can detect it? And so we'll talk about those two issues one after the other.

So first of all is the neutral production. And I guess we could call the second problem neutral transport. So in terms of neutral production, there are two major processes that we could have. We could have a charge exchange. So I'm just going to write that as CX. We talked about charge exchange before. This is where a neutral collides with an ion, and the neutral loses an electron and the ion gains an electron.

So we could have charge exchange with edge neutrals. And they could be neutrals coming in from the edge of the plasma. And that would look like some neutrals coming in like this, and I'll use n_a to be some density of neutral atoms. These are going to be wandering in from the edge all the time. Maybe a few of them make it to the center, and then they charge exchange of our core. And then those core-- [INAUDIBLE]. Hence, the ions make it back out again and we can detect them. That's one possible mechanism.

A more sort of active mechanism would be taking a neutral beam and putting a neutral beam into the center of the core here. This is neutral beam injection. We use neutral beam injection for all sorts of things, like current drive and heating of the plasma. And so we often have this process happening automatically.

The whole point of the neutral beam here is that as these neutral particles come into the center, they'll charge exchange with the ions and become ions. But that means that the ions they charge exchange with become neutrals. And so those neutrals may come out of the plasma here, where we can detect them. So we have charge exchange with edge neutrals and charge exchange with neutral beam injection.

Neutral transport is asking, can the neutral make it from the center of the core-- yeah. Can that neutral make it out along some path? And so there's going to be a series of processes which are trying to re-ionize this neutral as it tries to leave the plasma.

So these processes are going to be things like-- I'm going to give them their cross-sections, σ_e . This is electron impact ionization, which we talked about before. This is going to be-- we can use σ_p . This is ion impact ionization, because we can also have collisions of ions causing ionization here.

And finally, we can have the probability of charge exchange with ions. So if we've charge exchanged once in order to get our neutral in the first place, it's not at all impossible for us to charge exchange again as we head out, and then our core neutral that was heading out and giving us information about the core will become an ion and it'll be trapped again, and now, we'll have an ion and further out become a neutral, and that will try and leak.

So this charge exchange process is particularly problematic because what we're ending up with is something that looks like a hydrogen neutral plus a hydrogen ion going to a hydrogen ion but a hydrogen neutral. Now, the beautiful symmetry in this equation here means that this is resonant. And so it's quite likely to happen. So this σ_C is large.

So that's an overview of what we're going to try and do. I'm going to talk a little bit about the transport, estimate the size of these different coefficients, and work out how likely it is we can actually get any neutrals from the core back out again. And then we'll talk about the neutral productions. I'm going to go 2 and then 1. But that's an overview. If you have any questions before we get going, this is a good time to ask.

AUDIENCE: [INAUDIBLE].

JACK HARE: OK, yes. You should take a closer look at Hutchinson's book for a more formal definition. But quantum mechanically, this is nice and likely to happen because we're starting-- or ending up with something very similar to what we've started. And so this process is likely to occur. Yeah. OK.

So again, we've got some blob of plasma here. We've got some particle, which is born at a point A. And we want to know if we can get to a point B, like this. So this is the core, and B is where we have our detector. If you look at this, you say, this looks an awful lot like our radiation transport problem. And so we could say that there's a probability of the particle getting from A to B. Or this could be a power of particles as well if I multiply it by the particle energy per unit time.

And that's going to look like the exponential of the line integral from A to B of some opacity, in this case, more to do with scattering than it is to do with absorption, which varies along the path integrated-- along the path, the L, and then the minus sign. So again, this looks like radiation transport.

So we talked about three different processes over there which might be important. And the importance of those different processes will depend a lot on the distribution functions that we have here. So if this is V and this is f of V here, we'll have an ion distribution function. And that's just the ions making up the plasma. And that will be relatively small in velocity space here. So this is f of i.

And we'll also have an electron distribution function, which is much broader, and that's because the electrons are less massive. So for a similar temperature, they have higher velocities. And we'll also assume that our neutrals coming out here have some distribution that is relatively narrow and shifted like this. So this is at a loss T, the A like this for the neutrals.

And so we're assuming here, to make our calculations, that the ion velocities are all much less than these core neutral velocities. And this means that the i minus the A is about V_A . It's useful when we're working out these cross-sections that have the difference between the two velocity populations.

And then we can write down, that our attenuation coefficient alpha here is going to look like one upon V_A , the electron impact cross-section σ_e evaluated at V_A and e plus the ion impact ionization σ_p , V_i evaluated at V_A . And that's times n_i plus the charge exchange cross-section, σ_C V_i evaluated at V_A times n_i , like that.

And we'll also now assume that our beam velocity is much less than-- or, sorry, the neutral atom velocity-- the neutral is trying to get out of the plasma-- is much less than the sort of average thermal velocity of the electrons so that we can have this nice ordering where V_i is much less than V_A much less than V_e .

And that allows us to further simplify this. We'll simplify it in two ways. The first one is that from the point of view of the electrons, the ion-- sorry, the neutrals are basically not moving. So we can set this to 0. And from the point of view of the neutrals, this V_i is approximately constant. It's a small value, and it's approximately constant. And so we can just take this V_i out. Sorry.

We can evaluate this V_i as just being V_A and we can take it outside, where it will nicely cancel with this V_A . So this gives us an expression that looks like $1 \text{ upon } V_A \sigma_e V_e$ evaluated for effectively 0 neutral velocity n_e plus just these two cross-sections, σ_p and σ_C , times n_i .

And we see that this term goes as $1 \text{ upon } V_A$ and this term doesn't. So if we're dealing with energetic enough neutrals coming out of the core, then we can imagine that this term could be relatively small. It will depend exactly on what σ_e , σ_p , and σ_C are, but we can imagine, for reasonable values, this term will cancel. So we end up with an expression that just has this. So this again, is, just roughly α . Neglecting this for large-- OK.

So then, you can put this back up into our probability of the particle getting through, and we just get exponential of minus $\sigma_p \text{ plus } \sigma_C \int n_i dl$, like this. And you can go into Hutchinson's book and see graphs of σ_p and σ_C , and σ_e , as well, for different plasma conditions.

And so if we just use an example that maybe we have 10 keV neutrals streaming out of the core, because that's the sort of temperature we'd like in a fusion reactor, so 10 keV neutrals, we've got a density of around 10 to the 20 per meter cubed. So this is on the dense side for an MCF reactor, but that's the sort of density we might need for reactable conditions.

Then, using this, we will get a length scale in this exponent, which is effectively the mean free path of the neutrals trying to leave the plasma. And that is going to be about 10 centimeters.

So for this relatively hot, relatively dense plasma, we're not going to see very many neutrals more than 10 centimeters away from the core. And probably, our reactor is going to have a diameter A of about 100 centimeters or so. So we're not going to see very many of these particles, which means that in order for these neutral diagnostics to work at all, we're going to require. But that require...

AUDIENCE: And-- Yes.

JACK HARE: Yeah.

AUDIENCE: Have a--

JACK HARE: OK, thank you. So we're going to require that the line integrated density here is going to be less than about 10 to the 19 per meter squared here. And if you do that, then you get λA more like a meter or so. So we need "low" density. I put it in quote marks because low means different things for different people. But we need a relatively low density tokamak or magnetic confinement fusion device in order to get neutrals escaping from the core. OK, questions on this? Yes, Audrey?

AUDIENCE: [INAUDIBLE]

JACK HARE: This is just, like, a feed of the neutrals coming out here. I realize, actually, drawing this, I don't think this is right at all. I don't think this is a delta function, because we're talking about a distribution. That's the whole thing we're trying to measure here. So I think that, actually, I should be drawing it like this and call it f of A .

So this is the distribution of the neutrals. It's something Maxwellian-ish. But the average velocity here, which I could call V_{Ta} , which is maybe here, is much larger than V_{Ti} , but much smaller than V_{Te} .

So this is just the idea that for these neutrals here, from the point of view of the neutrals, the ions are basically stationary, and from the point of view of the electrons, the neutrals are basically stationary as well. And that allows us to simplify all of these angle brackets, which are actually averaging over the difference between the two velocity distributions. We can even set it to 0, or just be the velocity of the neutrals. Yeah?

AUDIENCE: [INAUDIBLE]

JACK HARE: Yeah, well, we imagine there's some temperature distribution inside this. So we're talking about, these neutrals here are neutrals that are in the core of the plasma, so we hope they're hotter than the ions in the rest of the plasma.

AUDIENCE: OK. So that [INAUDIBLE].

JACK HARE: No, these are the ions out here or something like that, further out. Yeah, that's a good question. So if we have a sort of radius and temperature-type thing, we assume, for a good reactor, that it's peaked in some way. And so we're asking you about this, the f of i V_i at r equals 0 becomes r of neutrals V_{neutral} at r equals 0. But they're streaming through some f of i V_i at r greater than 0, which in general is going to be a lower temperature. So the distribution is smaller. Yeah.

All of this is very hand-wavy, of course. If you want to do this properly, you have to keep track and do all these integrals correctly. But just to give an idea of the most important effect, it turns out r -- the [? ion-ion ?] impact ionization and the charge exchange. And the charge exchange dominates this pretty strongly, actually. So it's mostly just charge exchange with some of the other ions. Yeah.

AUDIENCE: [INAUDIBLE]

JACK HARE: We haven't talked about that. No. So what I'm saying is, imagine somehow you have created neutrals in the core. Would they actually be able to get out? Because if they can't get out, there's no point even trying to create them in the first place. So we're sort of working our way backwards through the problem.

The next step is, OK, now that we know, if we create some neutrals, they can escape, we can probably measure them in some way and learn about the core ion distribution. Now that we know that that's the case, can we now work out how to create the neutrals? And we'll talk about some techniques for that. Yeah. Yeah.

AUDIENCE: Why would we [INAUDIBLE]?

JACK HARE: Yeah. I mean, this is very crude. You're quite right. I mean, this should be inside here. But if we assume that these don't vary too much as the particle travels from the core to the edge, then it allows us just to pull out this factor of $n_i dl$ and make it clear that it's like the line-integrated density that's important thing. In reality, it would have to stay inside. Yeah. Cool. Any questions online? All right. Let's go find out how to put some neutrals in the core.

I'm trying to spell neutrals with an A there. OK. So again, Let's have some flux surfaces, some plasma like this. So one of the ways the neutrals could get there, as we discussed before, is from the edge. So around the edge of the plasma, there's, for example, the scrape-off layer. And then there's going to be a region which is in contact with the wall, which is relatively cold. And so we can have neutrals being produced here.

And these neutrals will wander in and they won't see the magnetic field lines, and so maybe one of them will only get so far here before the charge exchanges. And this one will get a little bit further. But you could imagine that maybe a couple of these neutrals wander into the center, they charge exchange, and then we can start getting the new neutrals, the ones that represent the core ions, start moving back out again.

So one possibility is-- we'll call it edge fueling. And we know how to calculate the probability of that, because we just spent a long time trying to calculate the probability of the opposite process. So now, we know that it's basically PBA, which is the probability of going from the outside, B, into the core, which we called A in our previous diagram.

So we can estimate how likely it is that any of these particles get through. And the answer is that this actually does create a fair amount of neutrals in the core. Just because there are so many neutrals in the edge, some of them will get through. So this technique by itself will create a detectable flux of core neutrals coming back out again. So PBA, which is edge to core.

Another thing you might say. Well, perhaps, just by some atomic processes in the core here, we could get neutrals developing here. So we could have radiative recombination, because, after all, there's a load of electrons flying around. These electrons are up above the ionization energy, so they're free. This electron could be caught by the electric potential of the ion and become a bound electron. And then we would get out our recombination photon.

That process could happen. But of course, the other process we've already discussed is when we have an electron come in and collide with one of the inner shell electrons, and we end up with two electrons coming back out again. And this is electron impact ionization.

And there'll be a balance in the steady state between these two processes and that balance will look like the density of electrons, the density of ions, the reactivity or the cross-section for the radiative recombination, timesed by the electron velocity. And that will be balanced by the electron density, the density of neutrals here, σ_e for electron impact ionization, V_e , like that.

So that means that we can look at this and go, OK, what is the steady state density of neutrals? That's going to be n_A over n_i . These cancel, so it's just going to be the ratio of these reactivities, $\sigma_r V_e$ over $\sigma_e V_e$, like that. And at 1 keV, which is a reasonable temperature for you to have in your core, this is about 10 to the minus 8 .

So the point I'm trying to make here is although you might occasionally have some recombination process, because the other processes which lead to ionization, again, are so rapid, you quickly will re-ionize. And so if you have 10 to the 13 or 10 to the 20 ions per cubic meter, they'll only have 10 to the 12 neutrals. And that's not really enough to detect. So this is negligible.

So edge fueling. Yes, it could be a way of getting neutrals in the core. Radiative recombination is not a way. And you kind of know this because if this was an effective process, then you would end up with a large neutral fraction in the core of your tokamak, and that's not something that we generally think about and doesn't generally happen.

And then finally, there are beams. And I'll talk a little bit more about beams in a second, as an active version of the diagnostic, but I just want to talk about how we actually measure these neutrals coming out in this passive version where we just let edge fueling do the work.

So passive cx with edge neutrals. So we've got some tokamak vacuum chamber, some MCF vacuum chamber, with some plasma like this. We have all of the particles coming in from the edge, and then the charge exchanging with particles in the core. And they're producing, for example, a hydrogen neutral that's coming out from the core.

And we put that hydrogen neutral into something called a stripping cell. And it comes out as a hydrogen ion again, but a hydrogen ion with the same velocity, or effectively, the same distribution function, as the neutral coming in. The stripping cell has a very electropositive gas, so it wants to take off the electrons, and it will charge exchange without really changing the trajectory of this very much.

And then we have, for example, a large magnetic field in this region that bends the particles according to their velocity. And then we have a detector like this with bins in it, and the bins have different sensitivity to different velocities here. And so by counting out a number of particles in each bin, we get out that distribution function, like that.

And we hope that, during their travels from the core through our stripping cell onto our detector, the distribution function has not been changed very much. So we really are measuring the distribution function of the particles in the center there. So if we look at that on our detector, I plot it on a logarithmic scale.

The log of intensity, which is equal to log of particle number versus particle energy here, but not velocity but energy, which is $1/2 mv^2$, we see that we have a region that is approximately straight, and that has a slope of $1/T_i$, because our distribution function for Maxwellian is exponential of $-e/T_i$ like that. And then we might have some interesting features down at the bottom here, where it looks non-Maxwellian. And so this can be the tail. And this will correspond to fast ions.

And measuring fast ions is very important in magnetic confinement fusion. These fast ions might be due to things like rf heating, or they might be due to runaways being produced by inductively driven electric fields, or other cool physics. So it's nice to be able to measure those fast ions as well.

The main problem with this technique, of course, is it's line-integrated. If I have my detector set up like this, I'm collecting all the particles in this region that come down here. And so you don't really know exactly where those particles have come from. And so we have the same problem we have with all of our other line-integrated diagnostics.

But it's pretty free. We don't actually have to put any effort into making neutrals in the center there. We just let the fact that we have a crappy vacuum chamber do the work for us. So this is a nice way to do it if you don't have a neutral beam. Any questions on that before we put a neutral beam into the system? Mm-hmm.

AUDIENCE: [INAUDIBLE]

JACK HARE: Oh, you would just have a magnetic field here. And frankly, the magnetic field of your tokamak does a pretty good job of that. But yeah.

AUDIENCE: How does it depend on the material of the [INAUDIBLE]?

JACK HARE: How does what depend?

AUDIENCE: Like, how many does [INAUDIBLE]?

JACK HARE: Yeah. I mean, we tend to be thinking about-- we still tend to be thinking about these neutrals coming from the edges being hydrogen. So when we have the hydrogen plasma-- by hydrogen, I mean deuterium and tritium, our fuel. So when we have our fuel in contact with the wall, there will be a layer of low ionization state, partially neutral gas around the edge, just because it's in contact with the wall. And so most of this will be coming from.

I don't know what happens if you start looking at this and you start looking at all the neutrals as well.

Presumably, you now need many more coefficients to calculate. And probably, if you're doing it properly, you should take into account those. But I think, although impurities can dominate things like line radiation, this really depends on the ion number. So we're going to be thinking, probably, about-- the impurities probably don't have a higher number density than the fuel, or if you've done something very wrong, that's happened. So I think you're still going to be dominated by your fuel ion. Yeah, cool. Yes?

AUDIENCE: So if you start with the neutral at the edge and then it goes into the plasma for charge exchange, then the new neutral that leaks out was an ion in plasma, and then you're measuring that neutral that leaks out. How do you know that you're only measuring both [INAUDIBLE] ion and the neutrals at the [INAUDIBLE]?

JACK HARE: Yeah. So the question was, how do we know that none of our signal is due to the neutrals at the edge? Well, the neutrals at the edge are very cold. So their energy is very low. So we probably will have a spectra that looks like this or something, with some noise up here, or some spurious signal that's due to other neutrals in the system. So I'll just start my fit at some nice high energy, which can't be due to any of these cold ions-- or cold neutrals from the edge. Has to just be due to the core ions instead. Any other-- yeah.

AUDIENCE: [INAUDIBLE],

JACK HARE: Yeah. So I mean, that's what's very challenging about this here, because in fact, this will be, like, the sum of a load of lines which will be weighted by their density. So I think it is difficult to interpret the passive version of it. Of course, if you go high enough out in energy and you can still fit a bit of a straight line, you can be like, there is some part of my plasma-- because you'll have different slopes.

And there will only be one slope that still extends all the way out here, because each of these will eventually sort of drop off. So if you're like, aha, still out here, there's some region, then that's probably your core ions. So you might have a chance of identifying them. But in general, your signal will look more complicated, which leads me very neatly on to using a neutral beam to do this instead.

So this actually is very similar to some of those active spectroscopy diagnostics we talked about before. We've, again, got our vacuum chamber and our plasma like this, but now, we have a neutral beam going through the plasma.

And if we have our detector oriented perpendicularly to this, now, we know that the majority of the neutrals are being born by charge exchange between the neutral beam and this small core region here, which overlaps the field of view of our detector and the neutral beam here. And so now, we know very well, when these particles come out, that we are detecting f_{iVi} about 0 from the core here. So this allows us to actively probe that central region here. It's the localizers.

And this is often convenient because we often have neutral beams. As we said, these neutral beams are already being used for heating and for current drive. Of course, the neutral beam has the same problem that we found out when we were talking about how far these neutrals can penetrate. If you have a density that's high enough inside your machine, you're not going to be able to use a neutral beam to heat all of your current drive.

So neutral beams tend to be limited to lower-density machines, which is why we've never had much use for them on devices like CMOS, where the magnetic field and the density is very, very high. But for other devices, you may well have a neutral beam already there. So this is, again, kind of a freebie to have this as a diagnostic.

Let's talk about some other things you can do with a neutral beam. So there's a very nice technique which is called-- where's my green? There we go. Charge exchange spectroscopy. This technique has developed a number of acronyms in the literature, which are all identical in meaning.

So some people call it charge exchange recombination spectroscopy, and some people call it charge exchange recombination spectroscopy. So if you see an acronym that looks something like this, they're probably talking about the same diagnostic. These are not two separate diagnostics. They are the same thing here.

So the idea of charge exchange recombination spectroscopy is, within our plasma, we will have some neutrals. Sorry, we will have some ions which are impurity ions. Sorry. So impurity, like A. We'll just use the symbol A. I'll make it clear what it is in a moment. These impurities will be things like oxygen or carbon or boron.

And although these impurities have many more electrons than the hydrogen, they are, in the center of the machine, probably fully stripped. They've lost all of their electrons. So they no longer can produce any light, so we can't detect their presence in the machine. So we don't know where these impurities are. So the idea of charge exchange recombination spectroscopy is, we put our neutral beam through the machine, and we end up charge exchanging on these impurities.

So we have some neutral beam and some impurity in a state Z plus where it's ionized at times, and this charge exchanges and gives us a neutral beam ion plus an impurity in a state $Z - 1$ plus. And most importantly, it's likely that the impurity will end up in an excited state.

So then our impurity, which is in some excited state, it will spontaneously decay down and we'll get out a photon like that. And because this photon will have a characteristic energy that we're looking for with our spectrometer, we can say, aha, this is from an impurity. And now, the invisible impurities are visible again.

So there are a few things that we can measure from charge exchange recombination spectroscopy. We can measure the density of impurities from intensity. So you need some atomic models and you need an absolute calibration for your detector, but you can work out the density of these impurities. That's pretty useful. We can also measure V_i and V flow from the Doppler broadening and the Doppler shift, respectively.

So as we discussed in the spectroscopy a few weeks ago-- or a week ago-- when we were talking about different broadening mechanisms, these effectively can be the impurities which are acting as a tracer for the flow velocity. So our bulk motion of the entire plasma around the torus and for the temperature.

And again, the reason is that maybe you could rely on impurity that was very high Z and still had some electrons left. But if you've done a really good job of reducing all of the tungsten and nasty things in your machine and all you're left with oxygen, carbon, and boron, then you need to effectively give them back an electron so you can do spectroscopy on them.

And again, this is a nice active diagnostic, because you can localize the region where the particles are sitting in by putting your spectrometer perpendicular to the line of sight of your neutral beam. And again, you can localize it to that region here. So this gives you these two, and they're all localized measurements.

There can be problems with this. So one issue is, you get charge exchange with edge impurities. I know I said this measurement localizes this, but you could imagine that you could get light reflecting from here go back down into your detector. And you can also just have background, so something like Bremsstrahlung radiation, at this photon energy.

And because there's not so many impurities, we hope, and not so many of them are charged exchanges neutral beam, the signal could be relatively weak. So there's Bremsstrahlung at the photon energy of the charge exchange with impurities. That could be a significant amount. So we may not be able to see this signal very clearly. So a really neat trick that you can do is to modulate the beam.

So you take the intensity of your neutral beam and you just turn it on and off like that. And then when you look at the intensity at this specific wavelength that you're on the lookout for, you should see that there's maybe some background, but that background is being modulated by the neutral beam.

And then you can do tricks like phase-locked loops and all your other favorite Fourier transform techniques, and you can isolate out which part of this signal is actually due to charge exchange and which parts of the signal is due to BREMS. That's pretty neat. OK, questions on this? Yes.

So theoretically, we have just a single line like this. If all the particles were stationary and the temperature was 0, all the particles moving at the same speed. And we just have a line at ω_0 . If all the particles are moving in the same direction, this will be Doppler shifted, and that Doppler shift would be adopted into the flow velocity.

And if we have some broadening on top of that, because we've got non-zero temperature, some of the particles-- although they're all moving in one direction. Some of them are moving back a little bit and some are moving forwards. And then the width of this feature, full width of half maximum, is proportional to $\sqrt{V T_i}$, which is square root of temperature. So you measure the shift of the feature and the broadening of the feature. That gives you the velocity and the temperature-independent. Any other questions? Any questions online?

OK, next topic. There are yet more fun things you can do with a neutral beam. Who would have thought it? Next fun thing you can do with a neutral beam is something called beam emission spectroscopy, BES.

The idea here-- once again, vacuum chamber, plasma, neutral beam. As the beam goes through the plasma, remember, our beam is going to be hydrogen or deuterium or something like that. The neutrals inside here-- and these are all neutrals.

So I'll just put a little 0 to remind ourselves they're neutrals-- will have some glancing collisions with the electrons and ions-- not enough to ionize them. Some of the collisions will ionize them. But some of the collisions will be enough to excite these. And so we'll end up inside our beam with excited neutral hydrogen and excited neutral deuterium or something like that.

And we can then put our spectrometer down here or our detector down here, and we can look for the characteristic-- we actually have some drawings of this. No. OK. We can look for these characteristic decays here. And if you remember before, we talked about the Lyman-alpha and -beta, and the Balmer series, as well. So we know very well-defined what these energies of the photons coming out are going to be.

Now, the neat thing about this is the intensity of these lines is proportional to-- or is-- I'll make it stronger. It's directly proportional to the number of electrons in the plasma, because this is a collisional excitation process and we know that that is proportional to N_e , and then whatever the relevant collisional excitation cross-section is. OK. But the main thing is it's proportional to the density.

So the neat thing is if I now have density fluctuations-- so this is-- we've got some background density, but the density is fluctuating on top of that. So $N_e = N_{e0} + \delta N_e$. Then that's going to give me an intensity fluctuation as well.

So by carefully measuring the intensity of these lines-- and again, we can localize this relatively well because I have a specific field of view for my detector and I know where this beam emission is coming from. So I can localize it to a little region. I could have multiple detectors. And I could have another one looking here. I can measure density fluctuations in very precise regions here.

And density fluctuations are very important, especially in magnetic confinement fusion, because these fluctuations lead to transport-- turbulent transport. So we really, really want to be able to [INAUDIBLE]. So beam emission spectroscopy, I think this is the only slide I have on this. Yes. Well, I have another one that's related.

This is a very cool technique. The trouble is that what I just said here about just being proportional to density here, this is rather simple. It's not actually true in reality. So what we actually have to do is include all relevant processes. This is going back to spectroscopy again, where if you're not quite in a simple equilibrium, you need to include lots more processes.

And you also need to include the beams' attenuation, because of course, the point of launching a neutral beam into your system is not to hit the far wall of your vacuum chamber and melt it, but it's actually to charge exchange with the plasma and heat it and drive current. So this beam is going to be getting less and less intense as it goes through the plasma.

And so if you want to do all this properly, you also need to track the beam attenuation, because this density is also proportional to the number of beam particles as well-- the number density of beam particles. So we need to track the beam attenuation.

But if you can do that, this is potentially a very powerful diagnostic for measuring density fluctuations, which are otherwise quite challenging to measure. Questions on beam emission spectroscopy?

AUDIENCE: Professor?

JACK HARE: Yes.

AUDIENCE: So do you need to specialize or neutral beam for these kind of readings, or is this something you can do while using the neutral beam as a heating source already?

JACK HARE: Yes. The nice thing about all of these things is you tend to just be able to use your standard neutral beam. So some people do have specifically made diagnostic neutral beams. So on eta, they have several neutral beams specifically designed for heating, and they also have diagnostic neutral beams. So the heating beams are up at megaelectronvolts, and the diagnostic beam is at a lower energy. And I believe that lower energy is because it gives you better cross-sections for all of these processes. So I'm not sure.

An amusing example of that is that there was a diagnostic neutral beam developed for the National Compact Stellarator Experiment at Princeton, which was canceled. And that diagnostic neutral beam then went on to be the heating neutral beam for a much smaller device. So you can build a diagnostic neutral beam, but at the end of the day, it's still a neutral beam. It can do whatever you normally want. And this is going to be coming in at 300 keV or whatever energy you can generate. And all these processes are going to be happening.

And the fact that the beam slows down and converts into ions and does all the heating relies on all of the same atomic processes that we then use to observe beam emission spectroscopy and charge exchange and all that sort of stuff. So those processes have to happen for the beam to work. And now, you're just observing the obvious consequence of them occurring. Yeah, good question.

AUDIENCE: Great, thank you.

JACK HARE: Other questions? Now we're going to be done early. The final thing, as we discussed it the other day, is the motional stark effect, which I'd forgotten was to do with this. And in fact, I was looking through my notes today and I was like, oh, I actually wrote several pages of notes on the motional stark effect and then told people I had no idea what it was. So that just shows how good my memory is.

So now, you will have the pleasure of learning about and then instantly forgetting about the motional stark effect. No, it's actually really cool. This is often abbreviated as MSE.

So we already talked about the stark effect. And in the stark effect, we said that there was an electric field. And I'm going to continue writing electric with this weird curly E's because we're going to have some capital E's for energy later and don't want us to get confused. So this is the electric field, which shifts the energy levels of some atom.

So we could do this just in a laboratory with a strong electric field and a neutral gas. And we then considered what this does inside a plasma. So in a plasma, we have lots of electrons moving around, as we said. And those electrons will be colliding with our ions, so our atoms that have bound electrons. And as those electrons get close, they will be shifting the electric field, and therefore, the energy level in the vicinity that atom.

So we saw there that we got a change in the energy of the photons being emitted, which was proportional to the electric field. And that electric field was proportional to density for the $1/3$. And this was just an electric field 1 over r squared argument. This should be $2/3$. There we go. And this was from collisions.

But this whole picture here required a thermal plasma with lots of particles moving in all sorts of ways here. So this was with the distribution f of V V_e that was exponential of, I don't know, minus V squared over V_{thermal} squared. But in this case here, for a neutral beam, we have a very peaked distribution function.

So our distribution of neutral beam ions is roughly a delta function. So they're all at V_B . So we don't expect to see a broadening. What we would expect to see, maybe, is some sort of shift in a single direction here but where is the electric field coming up from in a plasma like this, where we're just talking about our particles moving?

Because here, we argued this was the electric field from particles close in within the Debye sphere. But on large scales, we don't see this electric field. But what we do have is a magnetic field. We have a magnetic field B , like this. And so we have some beam velocity V cross B . And in the frame of the particle, that creates an electric field, which is equal to minus V cross B . So this is our Lorentz transform.

So in the frame of the lab we see particles moving with velocity V_{beam} , and we see a magnetic field. But in the frame of the particle, when we do the Lorentz transformation, the particle, of course, doesn't know it's moving within its frame. And instead of seeing a magnetic field, it sees this electric field instead.

And so that means that now, we're expecting to get energy shifts ΔE , which are, again, proportional to the electric field, which are now proportional to V cross B . And I just want to point out-- and Hutchinson goes into this in some detail-- we've assumed here that we're talking about the linear Stark effect.

And depending on exactly what's going on you may need more complicated things like the quadratic Stark effect. But this is the linear Stark effect. It turns out the linear Stark effect is good when this ΔE is large. So as long as you have sufficiently large magnetic fields and beam velocities, then you end up with this linear regime, which is much easier to analyze here OK. I'm going to erase this because I want to keep that on the board.

Now, some voodoo magic happens, and we find out that each energy level splits into multiple energy levels. So this is reminiscent of the Zeeman splitting again. So each energy level splits into energy levels whose energy is given by $3nK$. I'll explain what all of these symbols mean in a moment. Electric fields over ze over $4\pi\epsilon_0 A_0$ squared. I haven't got, in my notes, what A_0 is and I've forgotten what it is off the top of my head.

So that's not very useful. Some constant is the most important thing. All this time the Rydberg energy, which you all remember is 13.6 electron volts. And this electric field is V cross B modulus of it. OK, so n here is the quantum number. n equals 1, 2, 3. And K , here for each quantum number n , can take multiple values, K equals 0 plus or minus 1, plus or minus 2 up to plus or minus n minus 1.

So for n equals 1, K can only take the value 0. At n equals 2, it can take three values. At n equals 3, it can take five values like this. And because these energy levels are all split-- and I'll draw some now so that you get some feel for what's going on. Let's see if I've got space to do it here.

We start out with n equals 1, n equals 2, n equals 3. This n equals 1 will just be split when n equals 1 level. The n equals 2 will be split to three levels. And the n equals 3, we split to five levels. And each of these levels is going to have an energy, ΔE , between them, which is given by this formula here.

And so by measuring this splitting, we can measure the magnetic field, because most of the stuff in here is just a constant. We should know our beam velocity because we're the ones injecting the neutral beam. So the only thing we don't know is the size of the magnetic field.

So we could try and use-- so again, this is still a neutral beam, so we've got one electron. So we could try and look at transitions from here down to here. And there should be a triplet of transitions all coming out with slightly different energies. What is the problem with using this? first of all, does anyone remember the name for this transition, and what's the problem with using it?

AUDIENCE: Is it a forbidden transition?

JACK HARE: It's not forbidden, but good question. Hydrogen. Transition from the first excited state down to the ground state is the Lyman-alpha. What's the wavelength of the Lyman-alpha? Yeah it's 121 nanometers.

So this is what's called vacuum ultraviolet. And the reason it's called vacuum ultraviolet is it only propagates in a vacuum. So it's a real pain to deal with. So we don't want to work with this transition here. Instead, what we want to do is use a transition from n equals 3 down to n equals 2. What's the name of this transition?

AUDIENCE: H-alpha BALMER

JACK HARE: Yeah, H-alpha BALMER. This is where the forbidden transitions that were mentioned come in here. So these have an angular momentum quantum number m of 210 minus 1 minus 2, 1, 0, minus 1 like this. And the selection rule for these that tells you whether this transition is forbidden or not is Δm is equal to-- have I got it written down here? Yeah. 0 and plus or minus 1.

So from any given level here, for this one, for example, the m equals 0 and the n equals 3, we can have a transition down to 0, down to minus 1, and down to 1. We get a triplet of lines from that. We'll also get another set of lines-- you can see this starts to get quite complicated quite quickly-- like this.

There is no transition down to minus 2 because it doesn't exist. And then finally, we'll have-- well, the minus 2, could guess, can go to minus 1. And the m equals 1 can go down to m equals 0. And m equals 1-- I'm running out of colors-- and the m equals 2 can go down to here. Sometimes.

So you see there's all sorts of different lines. But you notice that their spacings of these levels are the same as the spacings of these. No, it's not, because n is different. Yeah, you're going to get clusters of different lines at slightly different energies. So you need a very nice spectrometer to be able to distinguish between these lines. So as you increase the magnetic field, the energy gap gets bigger, and so you need a less and less nice spectrometer.

So your resolution requirements go down for a bigger magnetic field. But in general, you're going to see a forest of lines. You're going to see, instead of just your h-alpha line that you had originally, you're going to get-- I'm not going to try and draw this accurately-- a set of lines like this. And then you might get another line over up here, and you might get another line down here, and all sorts of things like that, and other sets of triplets and doublets different places.

If that's not quite right, it's not completed. OK, cool. Why are we bothering with this? This looks an awful lot like Zeeman splitting. Couldn't we have just done Zeeman splitting? So the main thing about this is this ΔE is much, much larger than we get for Zeeman.

By "much, much larger," I about 10 times larger. So it's much easier to distinguish between the lines. And once again, we can use our neutral beam that we already have to do this. So this is a very nice use of a neutral beam.

Now, we talked about, with Zeeman before, it's not actually super useful to be able to measure the strength of a magnetic field in a tokamak, because we know what it is. It's the toroidal magnetic field. Basically, the poloidal magnetic field is very small. So this is our condition that the poloidal is much, much less than the toroidal that we have in a tokamak.

And the nice thing about this is that each of these different transitions has a different polarization. And this is going back to what we were talking about before, where we had these high polarizations and sigma polarizations. And I'm running out of room to draw it. But if you go back to the notes that we had on Zeeman spectroscopy here, this polarization, for example, will come out.

We've got some magnetic fields here, which is the toroidal field plus the poloidal field. We'll have the pi polarization coming out of the plasma perpendicular to B. And we'll have the sigma polarizations coming out parallel to B. So by having two spectrometers with two different polarizers cross of each other and measuring the relative strength of these two lines, you can actually determine the angle of the magnetic field.

And then from that, you can work out the size of this B poloidal component. And from the B poloidal component, you can start to guess at distribution of current inside your tokamak, which is very, very important, especially when we're doing things like bootstrap current, where we need very specific current profiles in order to induce it. So basically, motional stark effect here is very useful for magnetic reconstruction. Questions? Yes.

AUDIENCE: Is [INAUDIBLE] always there [INAUDIBLE]?

JACK HARE: So the Zeeman effect requires certain electronic configurations to actually occur. And I believe that you don't get Zeeman splitting in hydrogen, normally. So this motional stark effect is therefore more useful in that case. Yeah. And this motional stark effect is particularly big because we're talking about very fast-moving neutral beam particles. So this isn't useful in a plasma-- in just looking at the thermal ions, because all these velocities are in different directions and they're relatively small.

So all that will give you is a broadening. The point is that this V is large and all in one direction. So you get a big splitting and it's consistent. Every single particle is emitting light with the same spectra. And so you can actually stand a chance of measuring. Any questions online? All right. You can have a quarter of your-- quarter of an hour of your lives back. Do whatever you want. And happy Thanksgiving.